

LEVEL

TECHNICAL REPORT PR 80-27-312

ADVISORY DECISION AIDS: A PROTOTYPE

DECISIONS AND DESIGNS INCORPORATED

John F. Patterson
L. Scott Randall
Richard R. Stewart

February 1981

AD A098640

MAY 8 1981

DTIC FILE COPY

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

**Command & Control
Decision & Forecasting
Systems Program**

DEFENSE SCIENCES OFFICE • CYBERNETICS TECHNOLOGY DIVISION
DEFENSE ADVANCED RESEARCH PROJECTS AGENCY

81 5 07 003

TECHNICAL REPORT PR 80-27-312

**ADVISORY DECISION AIDS:
A PROTOTYPE**

by

John F. Patterson, L. Scott Randall, and Richard R. Stewart

Sponsored by

Defense Advanced Research Projects Agency
Contract MDA903-80-C-0194
DARPA Order No. 3831

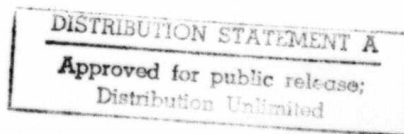
February 1981

THE VIEWS AND CONCLUSIONS CONTAINED IN THIS DOCUMENT ARE THOSE OF THE AUTHOR AND SHOULD NOT BE INTERPRETED AS NECESSARILY REPRESENTING THE OFFICIAL POLICIES, EITHER EXPRESSED OR IMPLIED, OF THE DEFENSE ADVANCED RESEARCH PROJECTS AGENCY OR THE UNITED STATES GOVERNMENT.



DECISIONS and DESIGNS, INC.

Suite 600, 8400 Westpark Drive
P.O. Box 907
McLean, Virginia 22101
(703) 821-2828



UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER PR-80-27-312	2. GOVT ACCESSION NO. AD-A098640	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) ADVISORY DECISION AIDS: A PROTOTYPE.	5. TYPE OF REPORT & PERIOD COVERED Technical Report	6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) John F. Patterson L. Scott/Randall Richard R. Stewart	8. CONTRACT OR GRANT NUMBER(s) MDA903-80-C-0194	9. PERFORMING ORGANIZATION NAME AND ADDRESS Decisions and Designs, Inc. Suite 600, 8400 Westpark Drive, P. O. Box 907 McLean, VA 22101
10. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency DSO/CTD, 1400 Wilson Boulevard Arlington, VA 22209	11. REPORT DATE Feb 1981	12. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
13. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)	14. NUMBER OF PAGES 43	15. SECURITY CLASS. (of this report) UNCLASSIFIED
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited		17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) DTIC ELECTE MAY 8 1981 C
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Decision aids Computer-aided analysis Mathematical models Decision analysis Problem solving Artificial intelligence Decision making Heuristic methods		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This report describes an effort to develop a prototype advisory decision aid (ADA). The three types of computerized decision aids used in problem-solving tasks can be described by three metaphors. One type of computer aid performs laborious tasks which the user would otherwise have to perform. Another type attempts to replace the user entirely and solves the problem alone. The third type of computer aid acts as an advisor. This decision aid has information about the problem and		

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

has procedures for making suggestions, but it does not solve the problem. That is left up to the user. As much as possible, it is desirable to place the functions of an advisor within computers. Such aids could be grafted onto data base management systems in an effort to help users cope with their information. In this capacity, the advisory aid is the natural extension of the HELP option being included in many operating systems. With time and continued improvements in our knowledge of artificial intelligence and computer-assisted decision making, advisory aids of a much more elaborate variety are likely to be available.

The initial ADA described in this report is not meant to be a fully operational aid; rather, its purpose is to clarify the nature and structure of ADAs. Fundamentally, the ADA concept seems quite sound. It requires much careful effort to construct the whole system, but in return one obtains an aid that can help a user work through his problem. For complicated problems requiring careful thought under stress, the detached approach of the ADA could prove quite beneficial. In any event, the ADA approach is probably the best vehicle for delivering heuristic techniques to the business, military, and government communities.

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	
Special	
A	

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

SUMMARY

This report describes an effort to develop a prototype Advisory Decision Aid (ADA). The problem domain is Air Force targeting for preplanned missions. The goal is to provide the Air Force analyst a tool which will guide him through an analysis of his problem.

One component of the ADA is a simulated environment. This is a program which permits the analyst to examine targets and any data that may be associated with them. In the present implementation this is accomplished using a videodisc mapping system. The mapping system retrieves maps from an optical videodisc and overlays symbols where the targets are located. These symbols can be colored based on target attributes.

While the simulated environment offers the user an unfiltered view of his problem, the remaining components of the ADA operate somewhat differently. The second component, the ADA's model of the environment, asks leading questions and makes suggestions concerning the user's problem. This program is a "devil's advocate" and promotes a more thorough examination of the problem by comparing its perspective of the environment with the user's. The purpose of this comparison is to construct the final component of the ADA: its model of the user's problem. Using heuristic techniques for drawing inferences from the data and the user's inputs, a cost/benefit model is constructed. This model identifies those targets which appear to provide the most benefit for the sorties that are available.

The prototype ADA is instructive because it suggests how artificial intelligence and decision analysis may be

blended to provide assistance to decision makers. Decision analysis is important for its ability to identify and formulate the user's perspective of a problem. Artificial Intelligence is important for its ability to define techniques for drawing inferences from data. Together the two should help to make data more accessible and useful to the user.

ACKNOWLEDGMENT

The authors wish to thank Audrey Wing Lipps and Dianne Waslov Caldwell for their superb implementation of the prototype Advisory Decision Aid (ADA) software.

CONTENTS

	<u>Page</u>
SUMMARY	iv
ACKNOWLEDGMENT	vi
FIGURES	viii
1.0 INTRODUCTION	1
1.1 A New Metaphor For a Computer Aid	1
1.2 The Development of a Prototype Advisory Decision Aid	2
2.0 THE PROBLEM DOMAIN	6
2.1 Target Nomination	6
2.2 Why Target Nomination?	7
3.0 THE SIMULATED ENVIRONMENT	10
3.1 The Implementation	11
3.1.1 The user interface	11
3.1.2 The mapping capabilities	13
3.1.3 The target database	14
3.2 Enhancements	15
4.0 THE ADVISORY AID	18
4.1 The Model of the Environment	18
4.1.1 The implementation	19
4.1.2 Enhancements	21
4.2 The Model of the User's Problem	24
4.2.1 The implementation	25
4.2.2 Enhancements	30
5.0 RECOMMENDATIONS AND CONCLUSIONS	32
REFERENCES	34
APPENDIX A: THE HARDWARE CONFIGURATION	A-1
APPENDIX B: LIST OF ORGANIZATIONS BRIEFED	B-1

FIGURES

		<u>Page</u>
Figure 1-1	THE THREE COMPONENTS OF AN ADA SYSTEM	4
Figure 3-1	USER'S PERSPECTIVE OF THE EQUIPMENT	12
Figure 4-1	REPRESENTATION OF AN INTERACTION IN THE DATA SPACE	23
Figure 4-2	GENERAL OUTLINE OF A COST/BENEFIT MODEL FOR TARGET NOMINATION	26
Figure 4-3	BENEFIT AS A FUNCTION OF NUMBER OF TARGETS FOR SYNERGISTIC AND REDUNDANT SETS	28
Figure A-1	HARDWARE CONFIGURATION	A-2

ADVISORY DECISION AIDS: A PROTOTYPE

1.0 INTRODUCTION

1.1 A New Metaphor For a Computer Aid

Computer aids to problem solving can be described in terms of two prevalent metaphors: Either they are tools or they are problem solvers. As tools, the aids perform laborious tasks which the user would otherwise have to perform. As problem solvers, the computer aid attempts to replace the user. The difference between the two types of aids derives from the user's attitude towards them. Either the aid is viewed as a means for extending the user's abilities or it is viewed as a means for replacing him.

There is also a third metaphor for a computer aid, that of the advisor. Here, the aid is neither an extension of the user nor an autonomous problem solver. Instead, it is another party to the problem-solving session. It has information about the problem and procedures for making suggestions, but it does not solve the problem. That is left to the user.

Under this metaphor, the aid is viewed as a knowledgeable subordinate. It reviews the problem solver's judgments and challenges him to think twice about decisions that seem unreasonable or unjustified. When the problem solver reaches an impasse, the aid attempts to move him towards a solution rather than allowing him to move in circles.

An advisory decision aid (ADA) of this sort offers the possibility of delivering a new type of computer assistance to the problem solver. Such an aid would employ heuristic computations, i.e., rules of thumb, rather than algorithmic ones. This would provide the ability to apply computers to

more complicated and more ill-defined problems, but only if the user will accept that the computer is capable of error. The aid's suggestions would be examined by the problem solver and rejected if they seem unreasonable or unwarranted. The goal is to provide an alternative perspective that the problem solver can match against his own. It is not required that this alternative perspective be infallible, merely that it be both reasonable and different.

This metaphor of the computer aid as advisor may seem overly anthropomorphic, but it serves a purpose. As much as possible, it is desirable to place the functions of an advisor within computers. Such aids could be grafted onto data base management systems in an effort to help users cope with their information. In this capacity, the advisory aid is the natural extension of the HELP option being included in many operating systems. With time and continued improvements in our knowledge of artificial intelligence and computer-assisted decision making, advisory aids of a much more elaborate variety are likely to be available.

1.2 The Development of a Prototype Advisory Decision Aid

This paper describes an effort to develop a prototype advisory decision aid (ADA). This initial ADA is not meant to be a fully operational aid; rather, its purpose is to clarify the nature and structure of ADAs. As a result, the prototype is quite austere and not yet fully implemented. It is, however, fully designed, and a discussion of this design should provide some insight into the nature of ADAs.

The first step in developing a prototype ADA is to identify a problem domain for the aid. Some concrete problem-solving situation is required both for clarity of exposition and as a technique to compel a realistic consideration of the difficulties inherent in constructing such an aid. It is our belief that a functioning prototype aid is a better

argument for the efficacy of the concept than abstract descriptions that have not yet faced the rigors of implementation. The problem domain that was chosen and the reasons for selecting it are discussed in Section 2.0.

The second step in the ADA development is to construct a computer simulation of the problem-solving environment. The term simulation is not intended to imply a frivolous or unrealistic environment, in fact, just the opposite. As much as possible, the simulated environment should attempt to mimic the situation that the problem solver would expect to encounter. Any information, tools, or other capabilities that the problem solver could expect to have at his disposal should be provided by the simulated environment. This is discussed in Section 3.0.

Figure 1-1 depicts the three major components of an ADA system. The user is provided access to both the ADA itself and the simulated environment. In the case of the environment, however, his interactions are monitored, but not modified, by the ADA. In addition, the ADA is also provided access to the simulated environment. Thus, the ADA system consists of two user interactive programs rather than one.

The simulated environment serves the same function as Winograd's (1972) block world. In his efforts to construct a language understanding program, Winograd found it necessary to build an artificial environment, i.e., the block world. This provided something about which he and his program could converse. An ADA confronts the same problem. The aid cannot fulfill its advisory function unless it has access to information about the problem. Otherwise, it has no basis for communication with the problem solver.

An ADA's simulated environment must, however, differ from the block world. Since the goal is effective problem

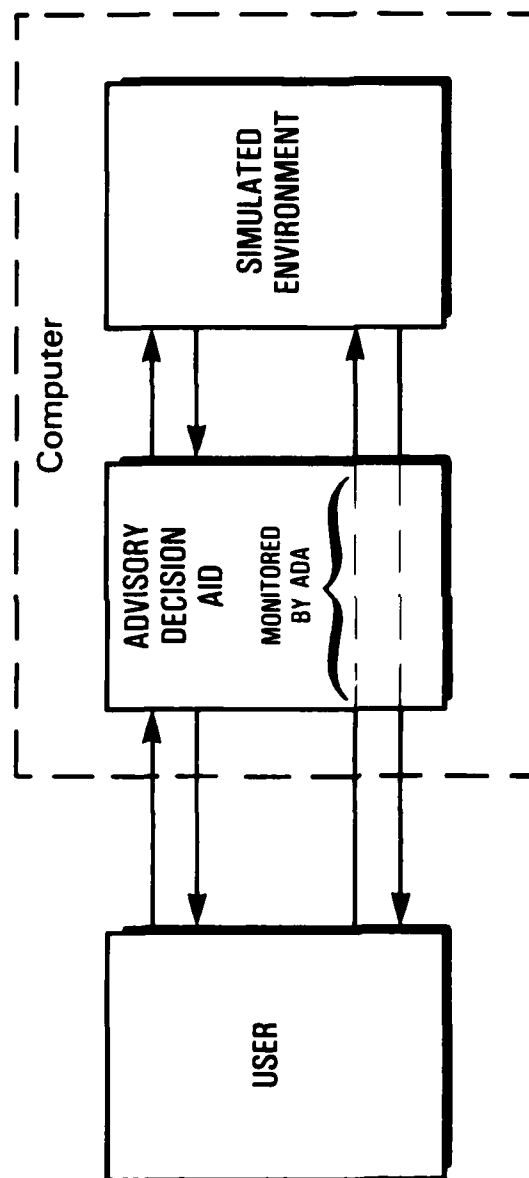


Figure 1-1
THE THREE COMPONENTS OF AN ADA SYSTEM

solving and not simply a demonstration of language understanding, the problem solver must be provided direct and unencumbered access to the simulated environment. In other words, the problem solver should not have to use the ADA to examine and manipulate the simulated environment. This is necessary because the problem solver may at some point need to challenge the ADA's advice by examining the situation directly. Thus, the simulated environment should allow the user to examine, manipulate, and otherwise investigate his problem even in the absence of an ADA.

The third step in the development of an ADA is the construction of the ADA itself. This consists of two components. The first is a model of the environment that permits the ADA to draw inferences about the nature of the problem. The second is a model of the user's problem, which only becomes complete through interaction with the problem solver. These are discussed in Section 4.0.

The distinction between the ADA's model of the environment and its model of the user's problem is an important one. Although the ADA may have a very sophisticated representation of the environment, this is not the same as knowledge of the problem. Only by learning about the problem solver's goals, intentions, and values can the ADA hope to construct a model of the problem. Moreover, it is quite likely that the problem solver will provide information to which the aid is not privy. Thus, the ADA can use its knowledge of the environment to challenge the problem solver, but it must also strive to develop a representation of the user's concerns.

The final step in the development of the prototype ADA is to evaluate its effectiveness and suggest how it can be improved. Discussions of this sort are provided in Sections 3.0 and 4.0, along with the discussions of implementation. In addition, Section 5.0 considers the conclusions and recommendations that can be derived from this effort.

2.0 THE PROBLEM DOMAIN

2.1 Target Nomination

The problem domain selected for the prototype ADA is Air Force target nomination. During a conflict, Air Force intelligence maintains information about a large number of enemy targets. This information is used to determine which targets provide the greatest opportunity for a bombing raid. Since there are many more targets than aircraft available to carry out the raids, it is important that targets be selected with skill and care.

Examination of an Air Force Targeting Primer (AFP200-17; 1978) reveals how difficult the problem can be. First, there are questions of the amount of benefit that can be expected by striking a particular target. This depends on whether the aircraft reaches the target, how much tonnage is dropped, how hard the target is, and how widely dispersed the target is. In addition, it depends on certain interactions between targets that may make it either unwise to strike one target after striking another, or unwise to stop after striking only one. There are also questions about the cost of a mission: What is the probability of aircraft loss or damage? What is the probability of a lost opportunity because all the aircraft are deployed? Finally, there are questions about the purpose, objectives, and strategy behind the bombing missions.

Despite these complexities, decisions must be made, and they must be well supported. According to the Targeting Primer,

"The results of detailed target analysis for preplanned mission targeting include an analysis model, a list of all installations comprising the target system, a chart with the installations plotted, photography and weapon-eering data on key and vulnerable facilities, and a

prioritized list of recommended targets."
(AFP200-17;1978).

The key product is, of course, the list of recommended targets. The other items largely serve to support the choice of certain targets over others.

Little would be served by further discussing the complexities of targeting. The prototype ADA does not capture all of the subtleties of the problem, and the reader can certainly appreciate the gist of the problem without understanding these details. Our purpose is not to educate the reader concerning targeting but to use targeting as a demonstration of ADAs.

2.2 Why Target Nomination?

Target nomination demonstrates four characteristics that make it especially well suited to an ADA. These are:

- 1) its importance and complexity;
- 2) the reasonableness of providing problem information via the computer;
- 3) the applicability of a well-defined analytic technique; and
- 4) the availability of a primitive problem element.

These characteristics are not especially restrictive, but they do clarify the situations in which an ADA may prove effective.

The importance and complexity of target nomination is probably the most obvious reason for its selection. Only an important and complex problem can justify the expense involved

in developing an ADA. Moreover, only a complicated problem can suitably challenge the concept of an ADA and thereby guide the development of that concept.

Also, target nomination is compatible with an ADA, because it is reasonable to suppose that the user's knowledge of the problem environment is provided by a computer. In this problem domain, the user is primarily examining data about targets, photographs of targets, and maps of the enemy territory, all of which can and are likely to be presented by computer. Thus, the ADA's need for a simulated environment is met without this seeming artificial or unnatural to the problem solver.

The third reason for preferring target nomination over other problem domains is that it conforms to the requirements of a well-defined analysis technique, namely, cost/benefit analysis. This permits the advisory aid to develop a meaningful representation of the user's problem. In the absence of this requirement, the aid could only react to the user's behavior and would have no ability to guide the user through the problem.

This is an important point and deserves close attention. In essence, the generic cost/benefit technique provides a framework or frame (Minsky, 1975) within which the advisory aid can represent the user's problem. Through its interactions with the user, the ADA must glean the information it needs to "fill the slots" of this generic model and thereby construct a specific model of the user's problem. The ADA's effort to complete the specific model provides a direction to the problem-solving session. Without this direction, the aid might lead the user in circles rather than converging on a solution.

The final reason for preferring target nomination as the problem domain is that it offers a primitive problem element out of which more complicated units can be defined. When the user and the advisory aid first interact, there must be some basis for communication. Although many of the concepts will be derived during the course of interaction, there must be some primitive or elementary concepts from which the more complicated concepts can be derived. In the case of target nomination, the individual targets are the primitive units.

Although these four reasons establish a preference for target nomination, there are many other problem domains that might have been chosen. In particular, any data base management problem in which the problem solver is perusing complicated data in an effort to reach a decision is a likely candidate for an ADA. The only requirement is for the user to accept that the ADA is an advisor and not the decision maker.

3.0 THE SIMULATED ENVIRONMENT

The simulated environment is the user's window on his problem. It provides him with the ability to examine and manipulate the data; it also offers him opportunities to see relationships between the problem elements. The simulated environment does not, however, interpret or otherwise filter the problem data. As much as possible, the data are provided in raw form and only transformed at the user's request. Presumably, this is what a good data base management system (DBMS) will do.

Assuming the DBMS is good, it is exactly what is needed for the simulated environment. One should not, however, adopt an overly restricted view of the nature of a DBMS or its capabilities. For one thing, the data may be of a type that is unfamiliar to standard DBMS efforts. It could consist of photographs, movies, maps, or sounds. Insofar as this is true, the simulated environment is more likely to resemble Negroponte's Spatial Data Management System (SDMS) (Negroponte, 1979) than other more conventional systems.

Also, the simulated environment should be more natural for the user than most DBMSs. He should spend as little time as possible on the formulation of queries and as much time as possible on viewing data and formulating his ideas. This is especially critical in an ADA system in which data are being perused or explored rather than simply accumulated according to certain predefined procedures. Natural modes of user interaction are also helpful to the ADA in that they decrease errors and thereby increase the likelihood that the user's queries actually reflect his underlying interests. This is important since the ADA is monitoring these queries for the purpose of understanding the user's interests.

3.1 The Implementation

Target nomination is an intrinsically geographic problem. The targets can be described in isolation, but their importance to and function for the enemy can only be understood in relation to geography. Not only are relative locations and distances important, but locations in relation to terrain, transportation systems, and other geographically distributed entities are equally critical. Thus, an important element of the simulated environment is an ability to display and manipulate maps.

In addition to the maps, there must also be information about the individual targets. Data of this sort are stored in a target data base. By overlaying information about the targets on the maps, an effective blending of the two types of information can be achieved.

Sections 3.1.2 and 3.1.3 describe these two types of data and the user's options for examining them. First, however, it is necessary to familiarize the reader with the equipment used by the ADA system. A full discussion of this equipment is presented in Appendix A. For the moment, only the user's perspective on the equipment need be discussed.

3.1.1 The user interface - From the user's point of view, the prototype ADA system consists of five components: a Ramtek color display system, an alphanumeric monitor, a keyboard, a joystick, and several special function keys (see Figure 3-1). The Ramtek color display is the principal component, the unit upon which maps are displayed. It is capable of reproducing a large number (256^3) of colors and color shades.

The alphanumeric monitor serves two purposes. First, it is linked to the keyboard and used to echo input

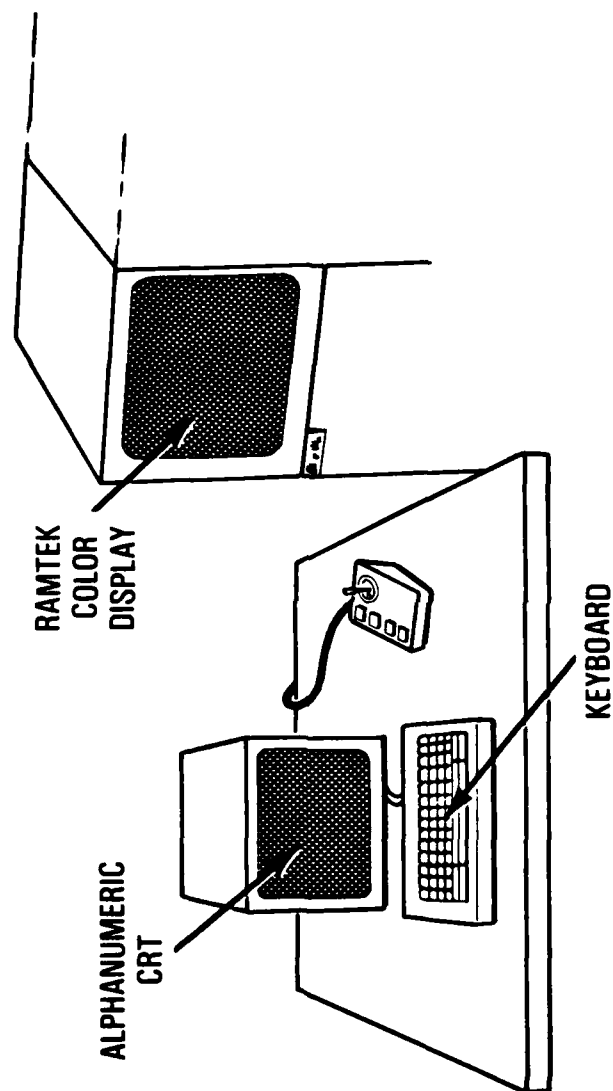


Figure 3-1
USER'S PERSPECTIVE OF THE EQUIPMENT

to the system from that device. Inputs of this sort are likely to be user requests and other control information. Second, the alphanumeric monitor is used to list data that the user requests.

The joystick and special function keys are housed within a single box and used to control the Ramtek Display. For example, the joystick can be used either to scroll across a map or to control the position of a cursor. The special function keys control such things as the level of zoom on a map and whether or not the cursor is activated.

3.1.2 The mapping capabilities - The setting selected for target nomination is North Korea. A standard 1:250,000 scale map of an area just north of the demilitarized zone was filmed and then used to master an optical videodisc. The take size for each frame on the videodisc was 3.2" x 2.4" in order to compensate for the loss of resolution introduced by the video medium. Thus, the nominal display scale on a 12" x 8" CRT would be about 1:62,500.

From the videodisc the maps can be retrieved and displayed on the Ramtek color monitor. This is done by first digitizing the videodisc output and placing the result in one of four frame buffers of a DeAnza Imagearray Processor, Model IP5532. The DeAnza is then used to control the Ramtek display.¹ (See Appendix A for details about the equipment.)

¹The Ramtek expects an RGB signal which can be provided by the DeAnza. The videodisc, however, represents color using the NTSC standard. This incompatibility was resolved by storing red, green, and blue separates of the map on the videodisc, reading each of these into a separate frame buffer, and driving the Ramtek's three channels from the three frame buffers. Interestingly, the overlay accuracy for the three separates was quite high (one pixel shift at most) despite the many opportunities for introducing error.

In the current implementation, the system retrieves a map segment from the videodisc. The user can zoom in (i.e., change scale) on this frame in powers of two and, assuming that he is zoomed in at least one level, he can scroll around (i.e., move laterally) on the image.

The final mapping capability currently available is the ability to overlay target symbols on the map. These are placed at their proper locations regardless of the user's level of zoom or choice of scroll.

3.1.3 The target database - The target database records information about targets in North Korea. The targets are currently limited to those that can be considered points on a map, but this is not a necessary restriction. For each of these targets, the database stores values for a variety of factors.

Three types of factor are used to describe a target. The first type identifies the name associated with a target, its location, the type of symbol used to represent the target, and so on. These factors are vital to the proper functioning of the system but are invisible to the user.

A second type of factor is of much greater concern to the user. These factors identify functional attributes of a target and thereby reflect its relation to the enemy's capabilities. For example, such factors identify whether a target serves an offensive or defensive purpose, whether it is a Navy, Air Force, or Army target, and whether it serves a C³, support, or combat role. The list of such attributes provides a functional description of the targets.

The final type of factor records information about the difficulty of destroying a target. At present, this is a single assessment of the number of sorties required to accomplish a specified level of damage. It could, however, be elaborated to include estimates of pilot risk, estimates of target hardness, or target dispersion, and estimates of a target's time to recover from an attack. These factors could easily be used to provide more accurate calculations of the cost of striking a target.

The user is provided two capabilities for examining these data. The first is by factor and varies the colors of the target symbols as a function of the designated factor. For example, if target function were selected, C^3 symbols would become green, support symbols would become orange, and combat symbols would become blue. For each factor that could be displayed, a specific color scheme was selected.

The second target data examination option is by target and provides a list of the data associated with a designated target. The desired target is indicated by moving a cursor onto its symbol and pressing a button. The requested data are listed on the alphanumeric monitor.

3.2 Enhancements

The simulated environment is, at present, rather austere. Map coverage is limited and the target data examination capabilities are simply queries by specific target or by factor. This is sufficient for a prototype, but there is much room for improvement.

In terms of the mapping capabilities, three enhancements are needed. The most important is to increase the coverage available to the user. As the user scrolls off of one frame of map data, the adjacent frame should be retrieved from the videodisc and displayed. Similarly, as the user zooms in on a map, he should be able to obtain maps with greater detail rather than simply viewing a larger version of the same map. Since the optical videodisc can store 54,000 video frames, the potential for large coverage is available. Indeed, this was the reason for choosing this method of map storage.

Another mapping enhancement is to provide map overlay control to the user. By storing separate maps of the roads, railroads, rivers, contour lines, etc., the computer can be used to construct a map that contains only those overlays that are relevant to the user's problem. As Anderson and Shapiro (1979) point out, this is a major advantage of computer cartography. Whereas standard hard copy maps must present all the information that any user is likely to desire, a computer map can be tailored to the application. This helps to decrease the clutter and increase the speed with which a user can comprehend the map's information.

A third mapping enhancement would be an ability to identify and highlight critical regions that are of interest to the user. For example, Air Force personnel would like to know how safe it is to fly over a particular location. This is a function of the location's distance from air defense units and could be displayed as shaded areas of vulnerability. Bright areas would be highly vulnerable and darker regions would be safe. Similarly, given a specified flight plan, a pilot might be interested in knowing the area on the ground from which he could be seen. Information of this sort is not stored with the maps or with the target data base, but must be computed using each.

In terms of the target data examination capabilities, three types of enhancement are needed: display enhancements, query enhancements, and data manipulation enhancements. Display enhancements are concerned with better ways to present information. Visual attributes such as symbol size, fade, and blink rate could be used besides color to represent underlying target data. In addition, these attributes might be usefully employed for the purpose of displaying more than one target variable at a time. For example, the blink rate of a target might indicate its importance, while the fade indicates the level of uncertainty about its position. Multidimensional displays of this sort are needed so that more information can be placed in front of the user.

Query and data manipulation enhancements improve the system's data base management capabilities. Besides target and factor queries, a user should be able to specify complex combinations of data attributes and have only those targets meeting these requirements displayed. Also, the user should be able to request that the system display computed values, such as means, minima, maxima, and rank orders. These capabilities are merely those that a good DBMS would have and, therefore, which an ADA system should have.

4.0 THE ADVISORY AID

As mentioned in the introduction, the ADA component of the overall system consists of two parts: a model of the environment and a model of the user's problem. The model of the environment captures the aid's view of how the problem elements relate to one another and how they create a problem. The model of the user's problem is built through interactions with the user; it captures his view of the problem. The models interact and provide information to one another, but they are somewhat different in emphasis.

4.1 The Model of the Environment

If the ADA is to fulfill its function, it must offer an alternative perspective of the problem environment. Using this alternative perspective, the aid can not only inform the problem solver, but it can help him to break away from any stereotyped or rigid approach that he may have adopted. It is not entirely necessary that this alternative perspective be absolutely correct or more comprehensive than the problem solver's; merely that it be suggestive of novel possibilities.

Actually, the model of the environment must achieve a balance between two competing needs. It must be sufficiently similar to the user's perspective to gain acceptance and sufficiently dissimilar to challenge his thinking. Therefore, in the presence of many users, the aid may need the ability to adjust the novelty of its perspective in order to accommodate to its different users. This capability must, however, await some future ADA.

4.1.1 The implementation - An advisor can offer suggestions by adopting either an analytic or an analogic approach. In the analytic approach, the problem solver's perspective is challenged on the basis of logical inconsistencies, absurd implications, or overlooked considerations. In other words, the advisor analyzes the problem solver's perspective as compared to his own in an effort to reconcile any differences of opinion and find the basis for the problem solver's reasoning.

The analogic approach is somewhat different. Instead of seeking logical inconsistencies, the advisor strives to find analogies or situations that are similar to those of interest to the user. This is done in an associative fashion rather than a deductive one. The suggested similar situation can then be compared to the problem solver's perspective in an effort to determine which aspects of the problem seem critical. Those aspects of the analogy that seem valid clarify the dimensions of the problem solver's concerns. Those aspects that seem invalid clarify the dimensions that are irrelevant to the problem solver's concerns.

This second or analogic approach is the one adopted for the prototype ADA. In particular, the ADA identifies targets that are analogous or similar to those that the user has already identified as interesting. Given a set of interesting targets, the ADA suggests new targets on the basis of their similarity to the identified set.

The procedure for calculating similarity is, at present, a very simple one. It uses the functional attributes in the target data base. Those targets that share a preponderance of characteristics are considered similar and those that differ on most variables are considered dissimilar.

In other words, a similarity metric is defined over this data space.²

Using the similarity metric, the ADA can suggest similar targets by the following procedure. First, the central tendency of the user's set of targets is calculated in order to determine a prototypical target that can represent the set. Then, all targets are assessed for their similarity to this prototypical target. Finally, any targets falling within a similarity threshold are offered as suggestions for the user's consideration. Using the same basic technique, the ADA can also identify those targets that the user has already included, but which are most dissimilar from the prototypical target. It can then suggest that these be dropped from consideration.

²The actual similarity function is of the following form:

$$S(i,j) = \sum_{\text{all } k} w_k d_k(i,j),$$

where S is the similarity, i and j are targets, k is indexed over data dimensions, w_k is the weight of dimension k , and $d_k(i,j)$ is the within dimension difference between the value associated with target i and the value associated with target j . This final difference measure must be flexible to take account of nominal, ordinal, and cardinal dimensions. In the case of nominal dimensions, the formula,

$$d_k(i,j) = 1 \text{ if } x_k(i) = x_k(j) \\ 0 \text{ if } x_k(i) \neq x_k(j)$$

where $x_k(i)$ is the value of target i on dimension k , seems the best approach. For cardinal dimensions, however, the formula,

$$d_k(i,j) = x_{k,\text{max}} - x_{k,\text{min}} - |x_k(i) - x_k(j)|$$

where $x_{k,\text{max}}$ and $x_{k,\text{min}}$ are the extreme values for the dimension, seems preferable.^k Ordinal dimensions are more problematic. First, they are scaled into cardinal dimensions, and then the appropriate difference equation is applied. The interested reader should consult Sneath and Sokal (1973) for a variety of other potential similarity functions.

The actual operation of this overall technique is as follows. First, the user identifies a target that seems interesting by placing the cursor on the target symbol and pressing a function key. Then, the ADA identifies similar targets. If one of these similar targets seems interesting, the user can add it to the set. If not, he can explicitly indicate that it should be dropped from consideration. After reviewing these first suggestions, he can request a new set of suggestions. Since this new set will be based on newly included targets, the prototype, and therefore the new suggestions, will shift away from their earlier form. By repeating this process, the user can build a set of targets that embody his concerns. If at any time this set seems absurdly large, it can be reduced by requesting suggestions of targets to drop.

This iterative process is used to help the user determine which targets deserve his attention. This target set is not the final list of nominated targets. In fact, it should contain two or three times as many targets as could be nominated. This target set simply reduces the scope of the user's considerations and allows him to concentrate a more detailed analysis on those targets that are prime candidates for nomination.

4.1.2 Enhancements - With the prototype ADA, the model of the environment is simply the data space plus a concept of similarity. Although the model is sufficient to demonstrate a concept, it is admittedly austere. To be truly effective, the ADA's model of the environment needs considerable improvement.

A major problem with the current implementation is that similarity is calculated over the data space. A more sophisticated approach would be first to map the data variables into a higher-level conceptual space and then to

calculate similarity in this derived space. This approach could allow very complicated mappings from data variables onto conceptual variables. These would be designed to eliminate discontinuities and interactions in the definition of concepts, thereby increasing the applicability of a linear similarity function.

Figure 4-1 suggests how a higher level space might be useful. For the two data variables, distance from the front and target function, there is an interaction in the definition of target importance. For logistics targets, the more important units are placed far from the front, while for the combat targets, the important units are near the front. This interaction renders the similarity function virtually meaningless, since it uses a linear calculation that would view target A as more similar to B or C than to D. By transferring the data space into the higher level variable, importance, this interaction can be overcome.

This question of interactions suggests another limitation on the ADA's model of the environment. It involves the manner in which the ADA draws inferences. There is little reason to believe that the user's targets of interest will cluster around a single prototypical target regardless of whether the data space or some conceptual space is used. Instead, it seems wiser to be prepared to calculate multiple prototypes, both for the targets of interest and for the targets that are explicitly excluded. Then, similarity could be based on the nearest prototype. This nearest prototype approach offers much greater flexibility to the ADA for capturing the user's concerns.

Finally, there is a question of whether the ADA should have a more sophisticated model of the enemy's capabilities and how its units interact. For example, a supply

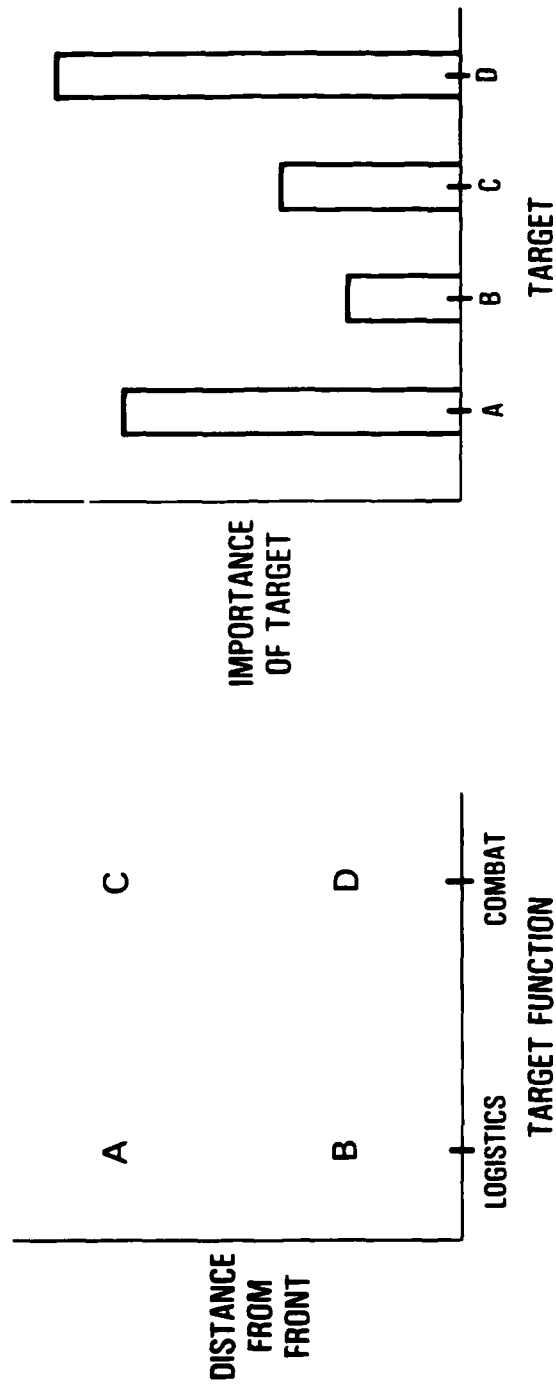


Figure 4-1
REPRESENTATION OF AN
INTERACTION IN THE DATA SPACE

network model could be useful for offering suggestions about the best place for an interdictive strike. Models of this sophistication are a favorite for operations research efforts and seem justified in the present context. Two issues must, however, be addressed: Will the data required be available? Will these models increase the ADA's rigidity so that its inferences are overly stereotyped?

Ultimately, the ADA's model of the environment should build upon the three disciplines of operations research, decision analysis, and artificial intelligence. From operations research, it can obtain techniques for modeling the physical world. From decision analysis, it can obtain techniques for capturing less well defined concerns, for representing these so that a user can understand them, and for combining this information into a single metric. And, from artificial intelligence, it can obtain procedures for drawing inferences from representations of knowledge. Achieving a blend of these three disciplines is perhaps the greatest challenge posed by the concept of an ADA.

4.2 The Model of the User's Problem

Besides its understanding of the environment, an ADA must also strive to represent and understand the user's problem. Although one might expect that a suitable model of the environment would permit the ADA to nominate targets without the user, this is not the case. The user has special information that he brings to the problem and about which the ADA is uninformed. He knows the stage of battle, the strategy being pursued, the condition of the troops, and any special constraints that might apply. In fact, the user is the one who brings all of the ill-defined characteristics of the problem to the problem-solving session.

Unfortunately, as important as these ill-defined characteristics are, the user is frequently unaware of them. He

will recognize when they are violated or ignored, but he may be unable to voice them in advance. As a result, the ADA must attempt to elicit the user's concerns by pushing him through the problem. The ADA's efforts to construct a model of the user's problem serve this goal of eliciting the user's concerns.

Of course, the other major purpose of this model is to provide support and rationale for the user's decision. A well-formulated model will collect and clarify the reasoning behind the user's choices and thereby help him to explain his decision. This can be especially important in a problem domain such as target nomination in which those who interact with the ADA are likely to be different from those who make or implement the decision.

4.2.1 The implementation - The ADA's model of the user's problem is a cost/benefit model. At the outset, the aid assumes that the user's problem will conform to the generic cost/benefit technique, and it attempts to prompt the user into defining the specifics of this model. For simplicity, the costs are assumed to be the sortie requirements of a target (which are stored in the data base), and benefit assessments are explicitly requested once the model is structured. Thus, the problem is reduced to one of structuring the model.

Figure 4-2 depicts the general outline of a cost/benefit model for target nomination. The 16 targets under consideration have been grouped into levels within target systems. Each level of each target system has been assessed for both cost (C) and benefit (B). The levels of each target system are mutually exclusive options, while the systems themselves are independent and compoundable options. In other words, a set of targets is nominated by choosing

	LEVEL 1	LEVEL 2	LEVEL 3	LEVEL 4	LEVEL 5
TARGET SYSTEM 1	NONE B - 0 C - 0	$T_1 + T_2$ B - 40 C - 15	LEVEL 2 + T_3 B - 50 C - 20	LEVEL 3 + $T_4 + T_5$ B - 55 C - 30	LEVEL 4 + T_6 B - 60 C - 50
TARGET SYSTEM 2	NONE B - 0 C - 0	$T_7 + T_8 + T_9$ B - 2 C - 1	LEVEL 2 + T_{10} B - 10 C - 5		
TARGET SYSTEM 3	NONE B - 0 C - 0	$T_{11} + T_{12}$ B - 100 C - 60			
TARGET SYSTEM 4	NONE B - 0 C - 0	T_{13} B - 25 C - 10	LEVEL 2 + T_{14} B - 40 C - 20	LEVEL 3 + $T_{15} + T_{16}$ B - 60 C - 40	

B - Benefit T_i = the i^{th} target
C - Cost

Figure 4-2
GENERAL OUTLINE OF A
COST/BENEFIT MODEL FOR TARGET NOMINATION

one level on each target system. The cost of this set is given by the sum of the costs for each level, and the benefit is given by the sum of the benefits.

The problem for the ADA is to help the user define this structure. The one principle that guides this effort is the fact that target systems must be independent. This is a requirement of the model which guarantees that adding costs and benefits over target systems will provide the proper total cost and benefit. The implication of this restriction is that any target interactions must be captured within a target system.

Again, for simplicity, it was assumed that costs (i.e., aircraft sorties) are always additive and that the only interactions between targets must be benefit interactions. As for benefit interactions, two types warrant consideration. One, called benefit synergy, occurs whenever the benefit of striking one target increases after another is struck. The other type of interaction, called benefit redundancy, occurs whenever the benefit of striking one target decreases after another target is struck.³

Figure 4-3 depicts the benefit of striking different quantities of targets within a synergistic or redundant set. An example of a synergistic set is several bridges crossing the same river. If one is eliminated, the enemy will simply shift to the others. Only as the friendly forces come close to eliminating all bridges will the full benefit be realized. An example of a redundant set is a

³ Experience has taught us that the terms synergy and redundancy are confusing. This confusion occurs, in part, because benefit redundancy from the friendly perspective is benefit synergy from the enemy perspective, and vice versa. We have not used new terms because none seems to avoid the problem while still capturing the basic ideas. The reader is, however, warned to think of benefit in terms of the friendly perspective.

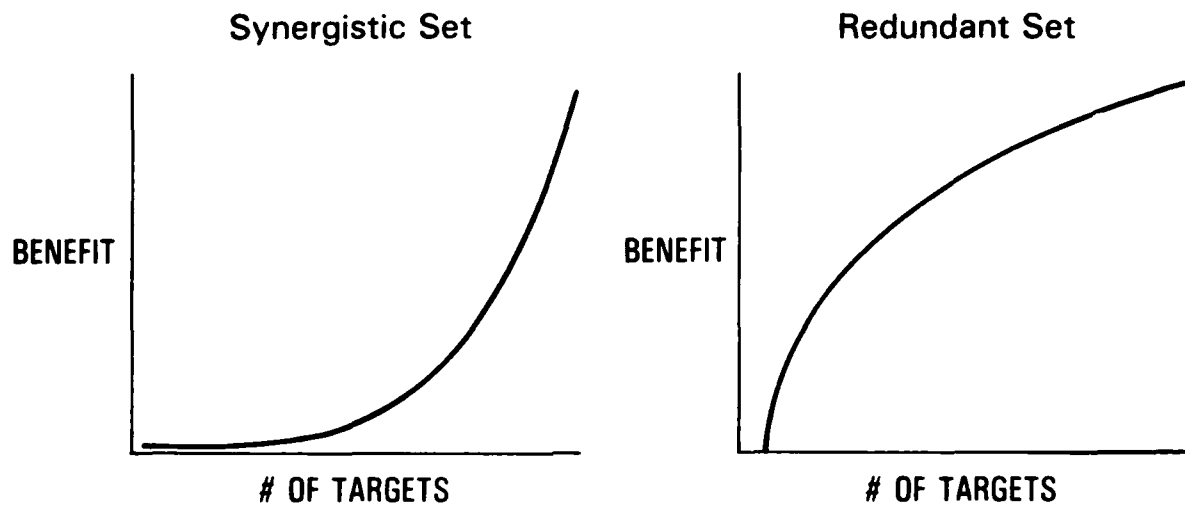


Figure 4-3
**BENEFIT AS A FUNCTION OF NUMBER OF
TARGETS FOR SYNERGISTIC AND REDUNDANT SETS**

series of bridges along a single road. If one is eliminated, most of the damage is done. Other strikes provide only diminishing returns.

The question remains, however, as to where these types of interactions can be captured by the cost/benefit model. If they must lie within the target systems, then they must be captured either within levels or between levels. This is, in fact, the policy that was adopted. Benefit synergy is to be captured within levels and benefit redundancy is to be captured over successive levels.

This is a major simplification and deserves close attention. Placing benefit synergy within levels implies that the synergistic set will be treated as all-or-none. Insofar as the benefit is zero until all are struck, this seems valid. However, the benefit curve (see Figure 4-3) is not likely to be so positively accelerated. As this curve approaches linearity, the accuracy of the model will deteriorate.

Another potential problem with this approach to capturing the interactions is that it assumes that synergistic sets bear a nested relationship to redundant sets. There is no a priori reason to believe this, and it has not been tested for the current problem. In the absence of justification, the approach must be applied cautiously and monitored for any inaccuracies that it may induce.

If this approach to capturing benefit interactions is accepted, then the next problem is to use it to induce a model structure from the user. The present effort approaches this in a very straightforward fashion. It is assumed that the user understands the distinction between synergistic and redundant interactions, and he is simply asked to group targets accordingly. First, he would identify targets for consideration by the techniques identified

earlier in Section 4.1.1. Then, he would partition these targets into mutually exclusive and exhaustive synergistic groups. Finally, these synergistic groups would be partitioned into mutually exclusive and exhaustive redundant groups. This would provide the structure, the costs would be in the data base, and the benefits would be explicitly elicited from the user. Thus, a model of the user's problem would be created.

4.2.2 Enhancements - The model of the user's environment is the aspect of the ADA system that is most speculative. The present design is by no means optimal and primarily serves to demonstrate the concept. Two major areas for improvement are the flexibility of the final representation and the procedures for eliciting the model.

As stated earlier, the present approach has a potential flaw, i.e., the assumption that target sets demonstrating benefit synergy are nested within target sets demonstrating benefit redundancy. Any model must, of course, have some simplifying assumptions, but it may be inadvisable to become so dependent on such an admittedly ad hoc one. The solution to this difficulty may lie in the development of several modeling approaches that can be selected as the need arises. This begs the question of how to choose among the approaches, but it may offer the only way to avoid the inappropriate assumptions of any one approach.

Another major problem with the current approach is that its procedures for eliciting information from the user are too tedious and blunt. It would be preferable to infer a structure from more natural inquiries such as: Are you interested in communications systems? Which region must be interdicted? This might require the ADA to draw more heavily on its model of the environment for the purpose of constructing a model of the user's problem, but this loss

of flexibility might be worth the improvement in the interactive quality of the aid. Indeed, the best solution might be to infer a model from some minimal set of questions and then adjust the model as the user disagrees with its implications. This would provide a secondary benefit in that a default model would be available in the event that the user were unable or unwilling to spend much time with the ADA.

Clearly, this problem of eliciting the model of the user's problem is one of the most difficult that an ADA implementation must confront. Not only is it necessary to specify adequate representations for the model, but it is also necessary to specify how the user's concerns will be captured. The present effort cannot claim to have solved these problems, merely to have pointed them out.

5.0 RECOMMENDATIONS AND CONCLUSIONS

At this time, the prototype ADA has been partially implemented and available for over six months. Several military and governmental groups have received demonstrations (see Appendix B), and their reactions have been generally favorable. These briefings are, however, of little use for evaluating the ADA concept, since they have emphasized the computer graphics rather than the advisory aid. Therefore, the reader must rely on the authors' opinions of the basic idea.

Fundamentally, the ADA concept seems quite sound. It requires much careful effort to construct the whole system, but in return one obtains an aid that can help a user work through his problem. For complicated problems requiring careful thought under stress, the detached approach of the ADA could prove quite beneficial. In any event, the ADA approach is probably the best vehicle for delivering heuristic techniques to the business, military, and governmental communities.

It is also fair to conclude that the current implementation is both austere and naive. In part, this is due to the desire to implement a prototype aid without delving too deeply into all the enhancements that suggest themselves. In addition, it is due to an initial lack of appreciation for the full dimensions of the problem. In any future effort, more time must be spent on the ADA's knowledge and less on the simulated environment.

Finally, an especially frustrating aspect of the current effort is that it has not been possible to conduct an evaluation of the ADA concept. Any future effort should be designed to move expeditiously towards a fully implemented ADA system

so that careful user evaluations can be conducted. These evaluations would serve both to question the merits of the concept and to offer guidance regarding the ADA's redesign or improvement.

On the whole, however, the prototype ADA has served its purpose. Its design has highlighted the issues that an ADA must confront, clarified the structure of this new type of aid, and suggested some ways in which the aid can help people reach decisions. Validation of the concept must be left to the future.

REFERENCES

- Anderson, R. H., and Shapiro, N. Z. Design Considerations for Computer-Based Interactive Map Display Systems. Final Report R-2382-ARPA. Santa Monica, CA: Rand Corp., February 1979.
- Minsky, M. "A Framework for Representing Knowledge." In The Psychology of Computer Vision, pp. 211-277. Edited by P. H. Winston. New York: McGraw-Hill, 1975.
- Negroponte, N. Spatial Data-Management. Cambridge, MA: Massachusetts Institute of Technology, 1979.
- Sneath, P. H., and Sokal, R. R. Numerical Taxonomy. San Francisco: W. H. Freeman and Company, 1973.
- U. S. Air Force. An Introduction to Air Force Targeting. AFPamphlet 200-17, October 1978.
- Winograd, T. Understanding Natural Language. New York: Academic Press, 1972.

APPENDIX A: THE HARDWARE CONFIGURATION

The hardware consists of two separate but interacting subsystems. One subsystem is designed to provide rapid access, manipulation, and display of whole image data (512 x 512 bytes). The second subsystem is a standard digital computer, which implements the non-image processing and acts as a host for the first subsystem. These interrelationships are clarified in Figure A-1 and in the sections that follow.

A.1 Whole Image Processing

The most distinctive aspect of the hardware is the whole image processing subsystem. This subsystem is virtually complete. Images that have been stored on a videodisc can be retrieved and placed in one of the DeAnza's four image memories. Here they can be scrolled, zoomed, added, and operated on in a variety of ways. Finally, these images can be displayed on the Ramtek color CRT. The primary capability lacking in this subsystem is an ability to write images onto peripheral storage.

Two additional components of the whole image processing subsystem are the videodisc controller and the time-based corrector. The videodisc controller is needed as an interface between the PDP 11/40 and the videodisc player. This provides the means to issue commands to the videodisc player.

The time-based corrector is needed to synchronize the output of the videodisc player with the DeAnza image processor. Because of minute signal fluctuations introduced by the videodisc player's servomechanism, its signal does not conform to the exacting requirements of the DeAnza's discrete timing. The time-based corrector compensates for these fluctuations.

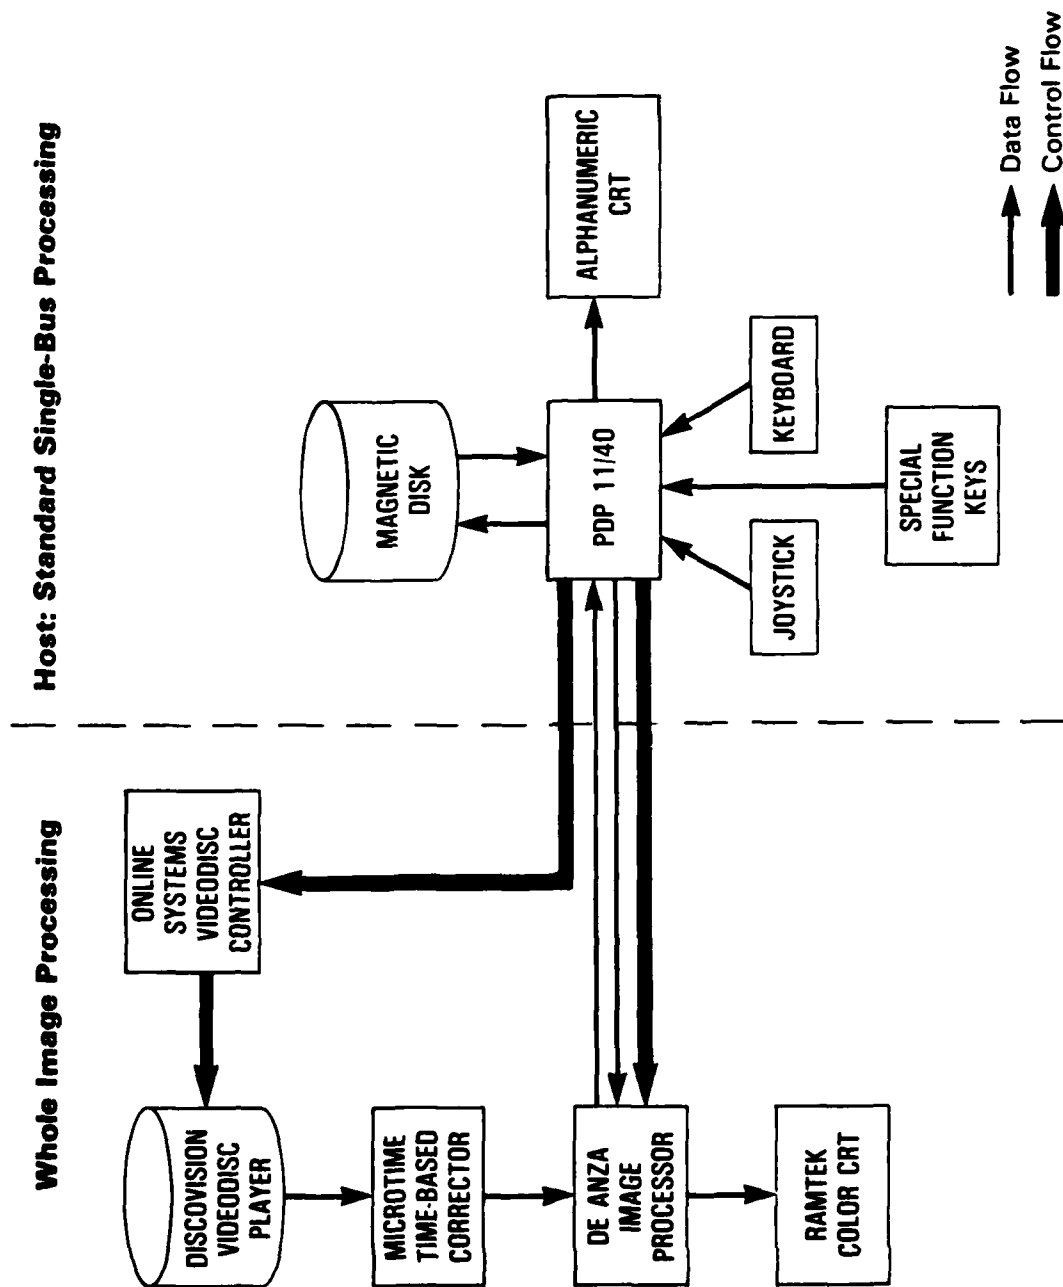


Figure A-1
HARDWARE CONFIGURATION

A.2 The Host

Besides operating on whole images, the system must also manipulate standard digital data. This is accomplished by a Digital Equipment Corporation PDP 11/40. This system has two RK05 magnetic disk drives for peripheral storage, as well as an alphanumeric CRT, a keyboard, and a DeAnza Systems joystick and special function keys for implementing the user interaction.

In addition to performing standard digital calculations, the PDP 11/40 serves as a controller for the image processor and videodisc controller. Thus, while the image processing and standard digital processing are implemented by separate devices, they are controlled by a single device. This arrangement not only provides a way to coordinate the various devices, but it also permits the use of well-developed programs, e.g., editors and compilers, when constructing the software to control the image processing devices.

A final feature of the hardware configuration is a facility to transmit information between the image processor memories and the PDP 11/40 memory. This capability allows portions of images to be modified as a function of information in the host's memory.

APPENDIX B: LIST OF ORGANIZATIONS BRIEFED

The following is a list of organizations that were briefed on the capabilities of the ADA system.

- Central Intelligence Agency
- Defense Intelligence Agency
- National Photographic Interpretation Center
- National Security Agency
- Office of Naval Research
- Pacific Data Systems Center
- Rome Air Development Center

CONTRACT DISTRIBUTION LIST
(Unclassified Technical Reports)

Director Advanced Research Projects Agency Attention: Program Management Office 1400 Wilson Boulevard Arlington, Virginia 22209	2 copies
Defense Technical Information Center Attention: DDC-TC Cameron Station Alexandria, Virginia 22314	12 copies
DCASMA Baltimore Office Attention: Mrs. Betty L. Driskill 300 East Joppa Road Towson, Maryland 21204	1 copy

SUPPLEMENTAL DISTRIBUTION LIST
(Unclassified Technical Reports)

Department of Defense

Director
Defense Advanced Research
Projects Agency
1400 Wilson Boulevard
Arlington, VA 22209

Director
Defense Advanced Research
Projects Agency
Defense Sciences Office
Cybernetics Technology Division
1400 Wilson Boulevard
Arlington, VA 22209

Chairman
Department of Curriculum Development
National War College
Ft. McNair, 4th and P Streets, SW
Washington, DC 20319

Defense Intelligence School
ATTN: Professor Douglas E. Hunter
Washington, DC 20374

Department of the Navy

Office of Naval Research
ATTN: Dr. Marty Tolcott, Code 455
800 North Quincy Street
Arlington, VA 22217

Office of Naval Research
ATTN: Dr. Bert King, Code 452
800 North Quincy Street
Arlington, VA 22217

Office of Naval Research
ATTN: Mr. Randy Simpson, Code 431
800 North Quincy Street
Arlington, VA 22217

Director
Naval Research Laboratory
Attention: Code 2627
Washington, DC 20375

Dr. Don Hirta
Naval War College
Newport, RI 02840

Dr. A. L. Slafkosky
Scientific Advisor
Commandant of the Marine Corps
(Code RD-1)
Washington, DC 20380

Chief
Intelligence Division
Marine Corps Development Center
Quantico, VA 22134

Dean of Research Administration
Naval Postgraduate School
ATTN: Patrick C. Parker
Monterey, CA 93940

Dean of the Academic Departments
U. S. Naval Academy
Annapolis, MD 21402

Dr. Glen R. Algaier
Head
Command & Control
Information Processing Branch
Naval Ocean Systems Center
San Diego, CA 92152

Department of the Army

Deputy Under Secretary of the Army
(Operations Research)
The Pentagon, Room 2E621
Washington, DC 20310

Director
Army Library - Army Studies (ASDIRS)
The Pentagon, Room 1A534
Washington, DC 20310

Dr. Edgar M. Johnson
Organizations and Systems
Research Laboratory
U. S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Technical Director
U. S. Army Concepts Analysis Agency
8120 Woodmont Avenue
Bethesda, MD 20014

Director
Strategic Studies Institute
U. S. Army Combat
Developments Command
Carlisle Barracks, PA 17013

Marine Corps Representative
U. S. Army War College
Carlisle Barracks, PA 17013

Commandant
Army Logistics Management Center
ATTN: DRXMC-LS-SCAD (ORSA)
Ft. Lee, VA 23801

Department of Engineering
ATTN: Col. A. F. Grum
United States Military Academy
West Point, NY 10996

Chief, Studies and Analysis Office
Headquarters, Army Training and
Doctrine Command
Ft. Monroe, VA 23351

Department of the Air Force

Assistant for Requirements
Development and Acquisition Programs
Office of the Deputy Chief of Staff
for Research and Development
The Pentagon, Room 4C331
Washington, DC 20330

Air Force Office of Scientific Research
Life Sciences Directorate
Building 410, Bolling AFB
Washington, DC 20332

Commandant
Air University
Maxwell AFB, AL 36112

Chief, Systems Effectiveness Branch
Human Engineering Division
ATTN: Dr. Donald A. Topmiller
Wright-Patterson AFB, OH 45433

Colonel Charles R. Margenthaler
Dean, School of Systems and Logistics
AFIT/LS
Wright-Patterson AFB, OH 45433

Commander, Rome Air Development Center
ATTN: Mr. John Atkinson
Griffis AFB
Rome, NY 13440

LTCOL Robert G. Gough
AFTEC/DAY
Curtland AFB, NM 87117

Other Government Agencies

Chief, Strategic Evaluation Center
Central Intelligence Agency
Headquarters, Room 2G24
Washington, DC 20505

Director
Center for the Study of Intelligence
Central Intelligence Agency
Attention: Mr. Dean Moor
Washington, DC 20505

Mr. Richard Heuer
27585 Via Sereno
Carmel, CA 93923

Other Institutions

Institute for Defense Analyses
Attention: Dr. Jesse Orlansky
400 Army Navy Drive
Arlington, VA 22202

Director, Social Science Research Institute
Attention: Dr. Ward Edwards
University of Southern California
Los Angeles, CA 90007

Perceptronics, Incorporated
Attention: Dr. Amos Freedy
6271 Variel Avenue
Woodland Hills, CA 91364

Stanford University
Attention: Dr. R. A. Howard
Stanford, CA 94305

Department of Psychology
Attention: Dr. Lawrence D. Phillips
Brunel University
Uxbridge, Middlesex UB8 3PH
ENGLAND

Decision Analysis Group
Attention: Dr. Miley W. Merkhofer
Stanford Research Institute
Menlo Park, CA 94025

Decision Research
1201 Oak Street
Eugene, OR 97401

Department of Psychology
Attention: Dr. Lee Roy Beach
University of Washington
Seattle, WA 98195

Department of Electrical and
Computer Engineering
Attention: Professor Kan Chen
University of Michigan
Ann Arbor, MI 94135

School of Engineering and
Applied Science
Attention: Dr. Andrew P. Sage
University of Virginia
Charlottesville, VA 22901

Harvard Business School
Attention: Professor Howard Raiffa
Morgan 302
Harvard University
Cambridge, MA 02163

Department of Psychology
Attention: Dr. Charles Gettys
Dale Hall Tower
University of Oklahoma
455 West Lindsey
Norman, OK 73069

Institute of Behavioral Science #3
Attention: Dr. Kenneth Hammond
Room 201
University of Colorado
Boulder, CO 80309